

PCAL: Language Support for Proof-Carrying Authorization Systems

Avik Chaudhuri[†] Deepak Garg

October 16, 2009
CMU-CS-09-141

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

[†]Author affiliation: University of Maryland, College Park

Avik Chaudhuri was supported by DARPA under grant no. ODOD.HR00110810073. Deepak Garg was supported partially by the iCAST project sponsored by the National Science Council, Taiwan, under grant no. NSC97-2745-P-001-001, and partially by the Air Force Research Laboratory under grant no. FA87500720028. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory, the DARPA, the U.S. Government, or Carnegie Mellon.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 16 OCT 2009		2. REPORT TYPE		3. DATES COVERED 00-00-2009 to 00-00-2009	
4. TITLE AND SUBTITLE PCAL: Language Support for Proof-Carrying Authorization Systems				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University ,School of Computer Science,Pittsburgh,PA,15213				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 33	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Keywords: Access control, proof-carrying authorization, language-based security

Abstract

By shifting the burden of proofs to the user, a proof-carrying authorization (PCA) system can automatically enforce complex access control policies. Unfortunately, managing those proofs can be a daunting task for the user. In this paper we develop a Bash-like language, PCAL, that can automate correct and efficient use of a PCA interface. Given a PCAL script, the PCAL compiler tries to statically construct the proofs required for executing the commands in the script, while re-using proofs to the extent possible and rewriting the script to construct the remaining proofs dynamically. We obtain a formal guarantee that if the policy does not change between compile time and run time, then the compiled script cannot fail due to access checks at run time.

1 Introduction

Proof-carrying authorization (PCA) [3, 5, 6, 18, 20, 22] is a modern access control technology, where an access control policy is formalized as a set of *logical formulas*, and a principal is allowed to perform an operation on a resource only if that principal can produce a *proof* showing that the policy *entails* that the principal may perform the operation on the resource. While this architecture allows automatic enforcement of complex access control policies, it substantially increases the burden of the user, since each request to perform an operation must be accompanied by one or more proofs. Furthermore, even if the user employs a theorem prover to construct the proofs, the user must still ensure that enough proofs are generated for each request to succeed, while minimizing the costs of proof construction at run time. In this paper we develop a programming language that can assist the user in performing such tasks correctly and automatically in a system with PCA. We have implemented a compiler for our language and tested it with a PCA-based file system, PCFS [18].

Our language, PCAL, extends the Bash scripting language with some PCA-specific annotations; the PCAL compiler translates programs with these annotations to ordinary Bash scripts, to be executed in a system with PCA. More precisely, PCAL annotations can specify what proofs the programmer expects to hold at particular program points. Based on these annotations, the compiler performs the following tasks.

1. It checks that the programmer’s expectations about proofs suffice to allow successful execution of every shell command in the script. For this, the compiler needs to know what permissions are required to execute each shell command. We provide this information through a configuration file.
2. Next, the compiler uses a theorem prover and information about the access control policy to try to *statically* construct proofs corresponding to the programmer’s annotations. In cases where static proof construction fails, because the annotations do not convey enough static information, the compiler generates code that constructs the proof at run time by calling the theorem prover from the command line.
3. Finally, the compiler adds code to pass appropriate proofs for each shell command to the PCA interface.

Thus, the output of the compiler is a Bash script which, beyond the usual commands, contains some code to generate proofs at run time (when it cannot generate such proofs at compile time), and some code to pass the proofs, generated either statically or dynamically, to the PCA interface.

Using PCAL offers at least two advantages over a naive approach, where a user generates and passes to the PCA interface enough proofs of access before running an unannotated script.

1. Because of the static checks and dynamic code generated by the compiler, it is guaranteed that the resulting script will *at least try to construct all necessary proofs of*

access. Thus, the script can fail only if the user does not have enough privileges to run it, and not because the user forgot to create some proofs. Indeed, we formally prove that if compilation of a program succeeds and the policy does not change between compilation and program execution, then the program cannot fail due to an access check (Theorem 4.2). This is very significant for scripts where the user cannot determine a priori what operations the script will perform.

2. Since the compiler sees all commands that the script will execute, it re-uses proofs to the extent possible and reduces the proof construction overhead, which a naive user may not be able to do. This is particularly relevant for POSIX-like policies where accessing a file requires an “execute” permission on all its ancestor directories. If several files in a directory need to be processed, there is no need to construct proofs for the ancestor directories again and again. The PCAL compiler takes advantage of this and other similar structure in policies and combines it with information about a program’s commands to minimize proof construction.

By design, PCAL and its compiler are largely independent of the logic used to express policies. The compiler requires a theorem prover compatible with the logic used, but it does not analyze formulas or proofs itself. Thus, the compiler can be (trivially) modified to use a different logic. Similarly, the compiler is parametric in the shell commands it supports. It assumes a map from each shell command to the permissions needed to execute it, and a single command to pass proofs to the PCA interface. By replacing this map and the command, the compiler can be used to support any PCA interface, not necessarily a file system.

PCAL is distinct from other work that combines PCA with a programming language [4, 20]. In all such prior work, the language is used to enforce access control statically. On the other hand, PCAL uses a *combination* of static checks and dynamic code to ensure compliance with the requirements of the PCA interface. Static enforcement is a special case of this approach, where an input program is rejected unless the compiler can construct all required proofs at compile time. Furthermore, in all prior work proofs are data or type structures and programmers must write explicit code to construct them. In particular, programmers must understand the logic. In contrast, PCAL separates proofs from programs, and shifts the burden of constructing proofs (and understanding the logic) from programmers to an automatic theorem prover. We believe that this not only makes PCAL’s design modular, but also easier to use.

Contributions

We believe that we are the first to propose, design, and implement a language that uses a combination of static checks and dynamic code to optimize the proof burden of a PCA-compliant program. This setting presents some unique technical challenges, and our design and implementation require some novel elements to deal with those challenges.

1. While we would like to discharge as many proofs as possible statically, we must be concerned about possibly invalidating the assumptions underlying those proofs at run

time. For instance, the state of the system may not remain invariant between compile time and run time. This requires a careful separation of (dynamic) state conditions from other (static) policies.

2. Since the values of some program variables cannot be determined at compile time, the PCAL compiler constructs *quantified proofs* which are parametric over these program variables. These variables are substituted at run time to obtain ground proofs. (See Section 5 for details.)
3. Programmer annotations in PCAL have both static and dynamic semantics. Statically, they *specify* authorization conditions and other constraints that should hold at run time, thereby aiding verification of correctness by the compiler. Dynamically, they *verify* any assumptions on the existence of authorization proofs and other constraints made by the compiler, thereby allowing sound optimizations.
4. For practical reasons, we must also be concerned about balancing the relative strengths and weaknesses of a theorem prover (to discharge proofs) and compiler (to analyze programs). We achieve such a balance by working at several levels of abstraction. While all functions and predicates used in a script have concrete implementations at run time, the compiler only partially interprets these functions and predicates with abstract rewrite rules, so that the program can be analyzed with appropriate precision by symbolic techniques. Furthermore, calls to the theorem prover are simplified, so that the theorem prover can treat all functions and predicates as uninterpreted, and thus can search for proofs efficiently. Tying these levels of abstraction together requires some care in the implementation. This is discussed further in Section 5.
5. We prove formally that the behavior of a compiled program is the same as that of the source program (Theorem 4.1) and that successfully compiled programs cannot fail due to access checks (Theorem 4.2). The proofs of these theorems require a precise characterization of assumptions on the theorem prover, the proof verifier, and the relation between the environment in which the program is compiled and that in which it is executed. We believe that this characterization is a significant contribution of this work, because it is fundamental to any architecture that uses a similar approach.

The rest of the paper is organized as follows. After closing this section with a brief review of related work, in Section 2 we discuss some background material covering PCA, and the assumptions we make about the interface it provides. Section 3 introduces PCAL and its compiler through an example. Details of the language, its compilation, and correctness theorems are covered in Section 4. Some important implementation-related issues are discussed in Section 5.

Related Work

There are two prior lines of work on combining proofs of authorization with languages. The first line of work includes the languages Aura [20, 24] and PCML₅ [4], where PCA as well

as a logic for expressing policies are embedded in the type system, and proofs are data or type structures that programs can analyze. This contrasts with PCAL, where proofs cannot be analyzed. PCAL’s approach is advantageous because it decouples the logic from the language, thus making it easy to use the same compiler with different logics. It also alleviates the programmer’s burden of understanding the logic. On the other hand, in Aura and PCML₅, parts of proofs can be re-used in different places, thus allowing potentially more efficient proof construction than in PCAL. However, it is unclear whether this advantage extends when automatic theorem provers are used in either Aura or PCML₅.

The second line of work includes several languages that culminate in the most recent F7 [8, 14]. These languages use an external logic like PCAL, but the objective is to express logical conditions. The programmer can introduce logical assumptions at different program points, and check statically at other program points that those assumptions entail some other formula(s). In PCAL it is not necessary that each programmer annotation about a proof succeed statically; if it fails, code to construct the proof at run time is automatically inserted. This approach is similar to hybrid typechecking [13], especially as applied to recent security type systems [9, 11]. Indeed, PCAL departs from previous lines of work in that it does not try to enforce security on its own; instead it is meant as a tool to help programs comply with a PCA interface that enforces security.

PCA, the architecture that PCAL supports, was introduced by Appel and Felten [3]. It has been applied in different settings including authorization for web services [5], the Grey system [6], and the file system PCFS [18]. The latter implementation is the basic test bench for PCAL. The specific logic used for writing policies in this paper (and PCFS) is BL [15, 16]. It is related to, but more expressive than, many other logics and languages for writing access policies (*e.g.*, [1, 2, 7, 12, 17, 19, 23]).

2 Background

In this section we provide a brief overview of PCA, and list particular assumptions that PCAL makes about the underlying PCA-based system interface.

PCA [3, 5, 6, 18, 20, 22] is a general architecture for enforcing access control in settings that require complex, rule-based policies. Policy rules are expressed as formulas in some fixed logic, and enforced automatically using formal proofs. Let \mathcal{L} denote a set of formulas that represent the access policy (see Section 3 for an example). The system interface grants user A permission η (*e.g.*, read, write) on a resource t (*e.g.*, a file) only if A produces a formal proof γ which shows that \mathcal{L} entails a formula **auth**(A, η, t) in the logic’s proof system. The formula **auth**(A, η, t) means that A has permission η on resource t . Its exact form depends on the logic in use and the resources being protected, but is irrelevant for the purposes of this paper. (Here it suffices to assume that **auth**(A, η, t) is an atomic formula.) The system interface checks the proof that A provides to make sure that it uses the logic’s inference rules correctly, and that it proves the intended formula. The system interface must provide a mechanism by which users can submit proofs either prior to or along with an access request. Even though users are free to construct proofs by any means they like,

it is convenient to have an automatic theorem prover to perform this task.

Assumptions

PCAL’s compiler supports rich logics for writing policies, in which proofs may depend not only on the formulas constituting the policy, but also on system state (*e.g.*, meta-data of files and clock time). Let H denote the system state. We write $\gamma :: H; \mathcal{L} \vdash s$ to mean that γ is a formal proof which shows that in the system state H , policies \mathcal{L} entail formula s . (In particular, s may be **auth**(A, η, t).)

PCAL assumes that an automatic theorem prover for the logic is available, both through an API and as a command line tool. A call to the theorem prover (either through the API or the command line) is formally summarized by the notation $H; \mathcal{L} \vdash s \searrow \gamma$, which means that asking the theorem prover to construct a proof for s from policy \mathcal{L} in state H results in the proof γ . Dually, $H; \mathcal{L} \vdash s \not\searrow$ means that the theorem prover fails to construct a corresponding proof. The latter *does not* imply the absence of a proof in the logic, since the theorem prover may implement an incomplete search procedure. The following command is assumed to invoke the prover from the command line and store in the file **pf** a proof which establishes **auth**(A, η, t) from the policies in **/pl** and the prevailing system state.

prove auth(A, η, t) **/pl** > **pf**

For passing proofs to the system interface, we assume a simple protocol: a command **inject** is called from the command line to give a proof to the system interface, which puts it in a store that is indexed by the triple (A, η, t) authorized by the proof. During the invocation of a system API, relevant proofs are retrieved from this store and checked. For example, the following command injects the proof in the file **pf** into the interface’s store.

inject pf

While proofs required to execute commands at run time must be ground, proofs produced at compile time may contain free variables, which we assume are listed in order, and which need to be instantiated at run time. For such proofs, the run-time substitutions of such variables are also provided to the **inject** command (with option **-subst**), so that the injected proofs are always ground. For example, the following command substitutes the run-time values of some Bash variables (in order) for the free variables in the proof read from file **.pf** and stores the resulting proof in the system interface.

inject pf -subst $\$_PRIN$ $\$z$ $\$x$ $\$y$ $\$bar$ $\$foo$

3 Overview of PCAL

In this section, we work through a small example to demonstrate the steps of our compilation. (PCAL is formalized in Section 4.) For this example, let there be a predicate **extension** and functions **path** and **base**, such that (informally):

- **extension**(f, e) holds if file f has extension e ;
- **path**(d, x) = p if path p is the concatenation of directory d and name x ;
- **base**(p) = x if **path**(d, x) = p for some directory d .

Consider the program P in Figure 1, written in PCAL. This program iterates through the files in some directory **foo** (unspecified), copying them to a directory **bar** (set to **"/tmp"**). Furthermore, it touches those files in **foo** that have extension **"log"**. The reader may ignore the **assert** statements (in lines 2, 8, 12, and 13) in a first reading; we explain their meaning below.

The system is configured to check, for any command, that certain permissions are held on certain paths in order to execute that command. Let us assume the following configuration:

Configuration

-
- Iterating over directory d requires permission **read** on d .
 - Executing the shell command **touch**(f) requires permission **write** on file f .
 - Executing the shell command **cp**(f_1, f_2) requires permission **read** on file f_1 , and permission **write** on file f_2 .
-

The **assert** statements in P serve to establish, at run time, that the principal running the script has particular permissions on particular paths. The compiler tries to statically identify **assert** statements that must succeed at run time, and eliminate them at compile time.

Assume that **member** is a predicate such that **member**(f, d) holds if file f is in directory d . Consider the following policy, written in a first-order logic with the convention that implication \Rightarrow is right associative.

Policy

```

auth("User", read, "/home").
 $\forall A. \forall x. \mathbf{auth}(A, \mathbf{write}, \mathbf{path}(\mathbf{"/tmp"}, x)).$ 
 $\forall A. \forall x. \forall y. \mathbf{member}(x, y) \Rightarrow \mathbf{auth}(A, \mathbf{read}, y) \Rightarrow$ 
    (auth( $A, \mathbf{read}, x$ )  $\wedge$ 
    (extension( $x, \mathbf{"log"}$ )  $\Rightarrow \mathbf{auth}(A, \mathbf{write}, x)$ )).

```

Informally, the policy asserts the following:

- the principal **"User"** has permission **read** on directory **"/home"**
- any principal A has permission **write** on any file in the directory **"/tmp"**

Program P

```
1 bar = "/tmp";
2 assert (read, foo);
3 for x in foo {
4   y = x;
5   x = base(x);
6   z = path(foo, x);
7   test extension(z, "log") {
8     assert (write, z);
9     shell touch(z)
10  };
11  z = path(bar, x);
12  assert (write, z);
13  assert (read, y);
14  shell cp(y, z)
15 }
```

Program Q

```
1 bar = "/tmp";
2 assert (read, foo);
3 for x in foo {
4   y = x;
5   x = base(x);
6   z = path(foo, x);
7   test extension(z, "log") {
8     -- assert (write, z);
9     shell touch(z)
10  };
11  z = path(bar, x);
12  -- assert (write, z);
13  -- assert (read, y);
14  shell cp(y, z)
15 }
```

Script S

```
#!/bin/bash
function base { _RET=${1##*/} }
function path { _RET=$1/$2 }
function extension { if [ ${1##*.} = $2 ]; then _RET="ok"; fi }
_PRIN="User"
foo="/home"

1 bar="/tmp"
2 prove auth ($_PRIN, read, $foo) /pl > pf
  inject pf
3 for x in `ls $foo`; do x=$foo/$x
4   y=$x
5   _RET="_"; base $x; x=$_RET
6   _RET="_"; path $foo $x; z=$_RET
7   _RET="_"; extension $z "log"; if [ $_RET = "ok" ]; then
8     inject /pf/1 -subst $_PRIN $z $x $y $bar $foo
9     touch $z
10  fi
11  _RET="_"; path $bar $x; z=$_RET
12  inject /pf/2 -subst $_PRIN $z $x $y $bar $foo
13  inject /pf/3 -subst $_PRIN $z $x $y $bar $foo
14  cp $y $z
15 done
```

Figure 1: Translation of an input program P , via an intermediate program Q , to an output script S . (The configuration, policy, and rewrite theory provided to the compiler are shown elsewhere.)

- for any principal A , file x , and directory y , if x is in y and A has permission **read** on y , then A has permission **read** on x , and furthermore, if x has extension "log" then A has permission **write** on x .

Finally, consider the following theory on the function symbols **path** and **base**, that abstracts the concrete semantics of these functions.

Theory

$$\forall x. \forall y. \text{member}(x, y) \Rightarrow \text{path}(y, \text{base}(x)) = x$$

Given the configuration, policy, and theory above, our compiler automatically translates P to the intermediate program Q in Figure 1. In Q , all **assert** statements except that in line 2 are eliminated, since the compiler can infer that they must succeed at run time. Such inference requires collection of path conditions, partial evaluation of terms modulo the given equational theory, and calls to the theorem prover. We describe the compiler in detail in Sections 4 and 5.

In particular, for the **assert** statement in line 8, the compiler reasons automatically as follows. Let **_PRIN** be the principal running the script. Line 8 is reached only if the following conditions hold for some z , x , x' , and **foo**:

- (1) **extension**(z , "log").
- (2) $z = \text{path}(\text{foo}, x)$.
- (3) $x = \text{base}(x')$.
- (4) **member**(x' , **foo**).
- (5) The statement **assert (read, foo)** in line 2 succeeds.

From condition (5), we can conclude that

- (6) **auth**(**_PRIN**, **read**, **foo**).

Simplifying conditions (2), (3), and (4) using the given theory, we have

- (7) $z = x'$.

Now from conditions (1), (4), (6), and (7) and the given policy, the theorem prover can conclude that **auth**(**_PRIN**, **write**, z), which is sufficient to eliminate the **assert** statement in line 8.

Next, we want to be able to run the intermediate program Q on a file system that supports PCA. The compiler translates Q to the equivalent Bash script \mathcal{S} in Figure 1. The commands **prove** and **inject** perform functions described in Section 2. The header (the part of \mathcal{S} before the numbered lines) defines all free variables (**_PRIN** and **foo**) and uninterpreted functions and predicates (**path**, **base**, **extension**) in P . The implementations

of such functions and predicates are sound with respect to the equational theory used by the compiler.

We close this section by discussing our trust assumptions. A policy is trusted, so any interpreted predicates in a policy (such as `member` and `extension`) must have trusted implementations (provided by the system). In contrast, a program is not trusted. The compiler may or may not be trusted. If the compiler is trusted, then the system can trust scripts produced by the compiler, and run such scripts without checking the proofs that they inject. This is significant in implementations where proofs may be large and proof verification may be costly. However, such a compiler cannot assume semantic properties of the functions used in a program (such as `base` and `path`) unless those functions have trusted implementations that are provided by the system. On the other hand, if the compiler is not trusted then the system must run all scripts with access checks. We implicitly assume the latter scenario in the sequel, and provide additional guarantees for the scenario in which the compiler is trusted (Theorem 4.2).

4 PCAL: Syntax, Semantics, and Compilation

We now describe the PCAL language and its compiler. We present the syntax of PCAL programs, define their operational semantics, formalize our compilation procedure and show that it preserves the behavior of programs.

For simplicity of presentation, we abstract various details of the implementation. (See Section 5 for a more detailed discussion.) Instead of Bash, we consider an extension of PCAL as the target language for compilation; programs in this target language can be easily rewritten to Bash. We also treat all function symbols as uninterpreted, although in principle, equations over terms may be freely added in the run time semantics (to model concrete implementations) and in the compiler (to model abstract properties of such implementations).

We assume that η , x , and t range over permissions, variables, and terms whose grammars are borrowed from the logic used to represent policies. φ denotes a logical predicate whose truth depends only on the system state (*i.e.*, a predicate that is not defined by logical rules). PCAL programs are sequences of statements e described by the grammar below. Directories, files, and paths are represented as terms, and χ is a special variable that is bound to the principal running a program.

Syntax

$e ::=$	statements
for x in t $\{P\}$	for each file f in directory t , bind x to f and do P
test φ $\{P\}$	if condition φ holds, do P
$x = t$	assign t to x
shell $n(t_1, \dots, t_k)$	call shell command n with parameters t_1, \dots, t_k
assert (η, t)	assert that principal χ has permission η on path t
$P, Q ::=$	programs

$e; Q$	run e , then do Q
end	skip/halt

We also consider below an extension of PCAL which acts as the target language for the compiler. $\alpha = \text{prove } (\eta, t)$ and $\text{inject } (\eta, t) \gamma$ are formal representations of the commands **prove** and **inject** from Section 2. γ ranges over proofs and α denotes a variable bound to a proof (which, in the actual implementation, is a temporary file that stores the proof; see Section 5).

Extended syntax

$e ::=$	statements
\dots	
$\alpha = \text{prove } (\eta, t)$	prove that principal χ has permission η on path t and bind the proof to α
$\text{inject } (\eta, t) \gamma$	inject proof γ that authorizes (χ, η, t)

Semantics

A PCAL program runs in an environment θ of the form (Δ, \mathcal{L}) , where Δ is a function from shell command names to lists of permissions (configuration) and \mathcal{L} is the set of logical formulas used to determine access (policy). Informally, if $\Delta(n) = \eta_1, \dots, \eta_k$ then executing shell command $n(t_1, \dots, t_k)$ requires permissions η_1, \dots, η_k on paths t_1, \dots, t_k respectively.

A state ρ is a triple (H, S, ξ) , where H is an abstract, logical representation of the part of the system state on which proofs of access depend, S is a function from paths to terms (data store), and ξ is a partial function from triples (A, η, t) to proofs (proof store). H must contain, at the least, information about members of directories. We write $\text{members}(H, t)$ to denote the list of files in directory t in the system state H . Proofs injected using $\text{inject } (\eta, t) \gamma$ are added to ξ .

Reductions are of the form $\rho, P \xrightarrow{\theta, \chi} \rho', P'$, meaning that program P at state ρ , run by principal χ in environment θ , reduces to program P' at state ρ' . The reduction rules are shown in Figure 2. $H, S \xrightarrow{n(t_1, \dots, t_k)} H', S'$ means that executing the shell command $n(t_1, \dots, t_k)$ updates the system state H and data store S to H' and S' respectively. $H \models \varphi$ means that φ holds in H , and $H \not\models \varphi$ means that φ does not hold in H . In practice, whether φ holds in H or not is decided using a trusted decision procedure provided by the system.

- **(Reduct for)** unrolls a loop P for each file x in a directory t . **(Reduct test)** simplifies $\text{test } \varphi \{P\}; Q$ to $P; Q$ if $H \models \varphi$, and to Q otherwise. **(Reduct assign)** is straightforward.
- **(Reduct shell)** finds proofs $\gamma_1, \dots, \gamma_n$ needed to authorize the shell command $n(t_1, \dots, t_k)$ in the proof store ξ . It then checks these proofs (premise $\gamma_i :: H; \mathcal{L} \vdash \text{auth}(\chi, \eta_i, t_i)$), and executes the shell command (premise $H, S \xrightarrow{n(t_1, \dots, t_k)} H', S'$).

Reduction $\rho, P \xrightarrow{\theta, \chi} \rho', P'$

(Reduct for)	$\frac{\rho = (H, -, -) \quad \text{members}(H, t) = t_1, \dots, t_k}{\rho, \text{for } x \text{ in } t \{P\}; Q \xrightarrow{\theta, \chi} \rho, P\{t_1/x\}; \dots; P\{t_k/x\}; Q}$
(Reduct test)	$\frac{\rho = (H, -, -) \quad H \models \varphi}{\rho, \text{test } \varphi \{P\}; Q \xrightarrow{\theta, \chi} \rho, P; Q} \quad \frac{\rho = (H, -, -) \quad H \not\models \varphi}{\rho, \text{test } \varphi \{P\}; Q \xrightarrow{\theta, \chi} \rho, Q}$
(Reduct assign)	$\rho, x = t; Q \xrightarrow{\theta, \chi} \rho, Q\{t/x\}$
(Reduct shell)	$\frac{\begin{array}{l} \theta = (\Delta, \mathcal{L}) \quad \Delta(n) = \eta_1, \dots, \eta_k \quad \rho = (H, S, \xi) \\ \xi(\chi, \eta_i, t_i) = \gamma_i \quad \gamma_i :: H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta_i, t_i) \\ H, S \xrightarrow{n(t_1, \dots, t_k)} H', S' \quad \rho' = (H', S', \xi) \end{array}}{\rho, \text{shell } n(t_1, \dots, t_k); P \xrightarrow{\theta, \chi} \rho', P}$
(Reduct assert)	$\frac{\begin{array}{l} \theta = (-, \mathcal{L}) \quad \rho = (H, S, \xi) \\ H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t) \searrow \gamma \quad \rho' = (H, S, \xi[(\chi, \eta, t) \mapsto \gamma]) \end{array}}{\rho, \text{assert } (\eta, t); P \xrightarrow{\theta, \chi} \rho', P}$
(Reduct prove)	$\frac{\theta = (-, \mathcal{L}) \quad \rho = (H, -, -) \quad H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t) \searrow \gamma}{\rho, \alpha = \text{prove } (\eta, t); P \xrightarrow{\theta, \chi} \rho, P\{\gamma/\alpha\}}$
(Reduct inject)	$\frac{\rho = (H, S, \xi) \quad \rho' = (H, S, \xi[(\chi, \eta, t) \mapsto \gamma])}{\rho, \text{inject } (\eta, t) \gamma; P \xrightarrow{\theta, \chi} \rho', P}$

Figure 2: Reduction rules

- **(Reduct assert)** calls the theorem prover to construct a proof γ which shows that χ has permission η on path t (premise $H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t) \searrow \gamma$), and passes it to the system interface by placing it in the store ξ .
- **(Reduct prove)** constructs a proof γ and binds α to it. **(Reduct inject)** places a proof γ in the proof store ξ . By these rules, the effect of the command sequence $\alpha = \text{prove } (\eta, t); \text{inject } (\eta, t) \alpha$ is exactly the same as the command $\text{assert } (\eta, t)$. However, $\text{assert } (\eta, t)$ occurs only in source programs whereas $\text{prove } (\eta, t)$ and $\text{inject } (\eta, t) \gamma$ occur only in compiled programs.

Compilation

Next, we formalize compilation of PCAL programs. As the compiler traverses a program, it maintains a database of facts that must be true at the program point that the compiler

is looking at. These facts are formally represented by $\Gamma = (\sigma, \Phi, \Xi)$.

- σ is a list of substitutions of the form $\{t/x\}$. The latter means that program variable x is bound to term t .
- Φ is a list of interpreted predicates φ that can be assumed to hold at a program point. These are gathered from commands **test** $\varphi \{ \dots \}$ and **for** x in $t \{ \dots \}$. In particular, φ may be of the form **member**(t', t), meaning that path t' is in directory t ; and we assume that **members**(H, t) = t_1, \dots, t_k implies $H \models \mathbf{member}(t_1, t) \wedge \dots \wedge \mathbf{member}(t_k, t)$.
- Ξ is a partial function from triples (A, η, t) to authorization proofs that the compiler has already constructed.

Figure 3 shows the rules to derive judgments of the form $\Gamma \vdash P \xrightarrow{H, \theta, \chi} P'$, meaning that under assumptions Γ , program P compiles to program P' in environment θ and system state H . χ is given to the compiler at the time of invocation; it represents the user who is expected to run the compiled program. H is the state of the system in which the compiled program is expected to run. It may either be the system state at the time of compilation (if it is expected that the compiled program will run in the same state), or it may be a state that the user provides. Both χ and H are needed to call the theorem prover during compilation.

For any syntactic entity \mathbb{E} , we write $\mathbb{E}\sigma$ to denote the result of applying the substitution σ to \mathbb{E} . $\mathcal{W}(P)$ denotes the variables that are assigned in the program P , and $\sigma \setminus \tilde{x}$ denotes the restriction of σ that removes the mappings for all variables in \tilde{x} . Finally, $|\Xi|$ and $\langle \Xi \rangle$ extract the formulas and proofs in Ξ (\prod denotes tupling of proofs):

$$|\Xi| = \bigwedge_{(A, \eta, t) \in \text{dom}(\Xi)} \mathbf{auth}(A, \eta, t) \quad \langle \Xi \rangle = \prod_{\gamma \in \text{rng}(\Xi)} \gamma$$

- **(Comp end)** terminates compilation when **end** is seen.
- **(Comp for)** compiles **for** x in $t \{ P \}; Q$ by compiling P to P' under the added assumption **member**($x, t\sigma$) (which must hold inside the body of the loop), and compiling Q to Q' . In each case, any prior substitutions for variables \tilde{x} assigned in P are removed from σ , because they may be invalidated during the execution of the loop (premises $\tilde{x} = \mathcal{W}(P/x)$ and $\sigma' = \sigma \setminus \tilde{x}$).
- **(Comp test)** is similar to **(Comp for)**; in this case the assumption $\varphi\sigma$ is added when compiling the body P of the test branch.
- **(Comp assign)** records the effect of assignment $x = t$ by augmenting substitution σ with $\{t\sigma/x\}$. This augmented substitution is used to compile the remaining program.
- **(Comp shell)** checks that there is a proof in the set of previously constructed proofs Ξ to authorize each permission needed to execute a shell command $n(t_1, \dots, t_k)$. (Proofs are added to this set in the next two rules).

Compilation $\Gamma \vdash P \xrightarrow{H, \theta, \chi} P'$

(Comp end)	$\Gamma \vdash \text{end} \xrightarrow{H, \theta, \chi} \text{end}$
(Comp for)	$\frac{\begin{array}{c} \Gamma = (\sigma, \Phi, \Xi) \quad \tilde{x} = \mathcal{W}(P) \quad \sigma' = \sigma \setminus \tilde{x} \\ x \text{ fresh in } \Gamma \quad \Phi' = \Phi, \text{member}(x, t\sigma) \\ (\sigma', \Phi', \Xi) \vdash P \xrightarrow{H, \theta, \chi} P' \quad (\sigma', \Phi, \Xi) \vdash Q \xrightarrow{H, \theta, \chi} Q' \end{array}}{\Gamma \vdash \text{for } x \text{ in } t \{P\}; Q \xrightarrow{H, \theta, \chi} \text{for } x \text{ in } t \{P'\}; Q'}$
(Comp test)	$\frac{\begin{array}{c} \Gamma = (\sigma, \Phi, \Xi) \quad \tilde{x} = \mathcal{W}(P) \quad \sigma' = \sigma \setminus \tilde{x} \quad \Phi' = \Phi, \varphi\sigma \\ (\sigma, \Phi', \Xi) \vdash P \xrightarrow{H, \theta, \chi} P' \quad (\sigma', \Phi, \Xi) \vdash Q \xrightarrow{H, \theta, \chi} Q' \end{array}}{\Gamma \vdash \text{test } \varphi \{P\}; Q \xrightarrow{H, \theta, \chi} \text{test } \varphi \{P'\}; Q'}$
(Comp assign)	$\frac{\Gamma = (\sigma, \Phi, \Xi) \quad \sigma' = \sigma[x \mapsto t\sigma] \quad (\sigma', \Phi, \Xi) \vdash P \xrightarrow{H, \theta, \chi} P'}{\Gamma \vdash x = t; P \xrightarrow{H, \theta, \chi} x = t; P'}$
(Comp shell)	$\frac{\begin{array}{c} \theta = (\Delta, _) \quad \Delta(n) = \eta_1, \dots, \eta_k \quad \Gamma = (\sigma, _, \Xi) \\ (\chi, \eta_i, t_i\sigma) \in \text{dom}(\Xi) \text{ for each } i \quad \Gamma \vdash P \xrightarrow{H, \theta, \chi} P' \end{array}}{\Gamma \vdash \text{shell } n(t_1, \dots, t_k); P \xrightarrow{H, \theta, \chi} \text{shell } n(t_1, \dots, t_k); P'}$
(Comp static)	$\frac{\begin{array}{c} \Gamma = (\sigma, \Phi, \Xi) \quad \theta = (_, \mathcal{L}) \\ H, \Phi; \mathcal{L} \vdash \Xi \Rightarrow \mathbf{auth}(\chi, \eta, t\sigma) \searrow \gamma' \quad \gamma = \gamma' \langle \Xi \rangle \\ \Xi' = \Xi[(\chi, \eta, t\sigma) \mapsto \gamma] \quad \Gamma' = (\sigma, \Phi, \Xi') \quad \Gamma' \vdash P \xrightarrow{H, \theta, \chi} P' \end{array}}{\Gamma \vdash \text{assert } (\eta, t); P \xrightarrow{H, \theta, \chi} \text{inject } (\eta, t) \gamma; P'}$
(Comp dynamic)	$\frac{\begin{array}{c} \Gamma = (\sigma, \Phi, \Xi) \quad \theta = (_, \mathcal{L}) \\ H, \Phi; \mathcal{L} \vdash \Xi \Rightarrow \mathbf{auth}(\chi, \eta, t\sigma) \searrow \alpha \text{ fresh in } \Gamma, P \\ \Xi' = \Xi[(\chi, \eta, t\sigma) \mapsto \alpha] \quad \Gamma' = (\sigma, \Phi, \Xi') \quad \Gamma' \vdash P \xrightarrow{H, \theta, \chi} P' \end{array}}{\Gamma \vdash \text{assert } (\eta, t); P \xrightarrow{H, \theta, \chi} \alpha = \text{prove } (\eta, t); \text{inject } (\eta, t) \alpha; P'}$

Figure 3: Compilation rules

- **(Comp static)** and **(Comp dynamic)** are used to compile the command **assert** (η, t) in different cases. To decide which rule to use, the compiler tries to statically prove $|\Xi| \Rightarrow \mathbf{auth}(\chi, \eta, t\sigma)$ by calling the theorem prover. The context in which the proof is constructed not only contains H and the policy \mathcal{L} , but also information about directory memberships and predicates tested in outer scopes $(\Phi\sigma)$. If a proof γ' can be constructed, rule **(Comp static)** is used: **assert** (η, t) is replaced by **inject** $(\eta, t) \gamma$, which passes the statically generated proof $\gamma = \gamma' \langle \Xi \rangle$ to the system interface at run

time. $(\gamma' \langle \Xi \rangle)$ is the proof of **auth** $(\chi, \eta, t\sigma)$ obtained by eliminating the connective \Rightarrow from $|\Xi| \Rightarrow \mathbf{auth}(\chi, \eta, t\sigma)$. Also, the fact that the new proof exists is recorded by updating Ξ to $\Xi' = \Xi[(\chi, \eta, t\sigma) \mapsto \gamma]$, and using Ξ' to compile the remaining program P . If the proof construction fails, rule **(Comp dynamic)** is used: the compiler generates code both to construct the proof at run time and to inject it into the system interface. Accordingly, **assert** (η, t) is compiled to $\alpha = \mathbf{prove}(\eta, t); \mathbf{inject}(\eta, t) \alpha$. Even in this case, it is safe to assume that a proof of **auth** $(\chi, \eta, t\sigma)$ will exist when P executes (else $\alpha = \mathbf{prove}(\eta, t)$ will block at run time), so Ξ is updated to $\Xi' = \Xi[(\chi, \eta, t\sigma) \mapsto \alpha]$.

Formal Guarantees

We close this section by stating two theorems that guarantee correctness of compilation. Proofs of these theorems can be found in the related technical report [10]. We begin by defining a preorder \leq on system states. Roughly, $H \leq H'$ if any formula that holds under H also holds under H' .

Definition 4.1 (\leq). *For any H and H' , let $H \leq H'$ if for all $\varphi, \gamma, \mathcal{L}$, and s , (1) $H \models \varphi$ implies $H' \models \varphi$, and (2) $\gamma :: H; \mathcal{L} \vdash s$ implies $\gamma :: H'; \mathcal{L} \vdash s$.*

Next, we assume the following axioms for the various external judgments. Roughly, Axiom (1) states that system states are updated monotonically by shell command executions. Axioms (2), (3), and (4) state that verification of proofs must be closed under substitution, modus ponens, and product. Axiom (5) states that the theorem prover produces only verifiable proofs (*i.e.*, the theorem prover is sound). Axiom (6) states that the theorem prover always produces a proof if some proof exists (*i.e.*, the theorem prover is complete).

Axioms

-
- (1) if $H, S \xrightarrow{n(t_1\sigma, \dots, t_k\sigma)} H', S'$ then $H \leq H'$
 - (2) if $\gamma :: H; \mathcal{L} \vdash s$ then $\gamma\sigma :: H\sigma; \mathcal{L} \vdash s\sigma$
 - (3) if $\gamma :: H; \mathcal{L} \vdash s$ and $\gamma' :: H; \mathcal{L} \vdash s \Rightarrow s'$ then $(\gamma' \gamma) :: H; \mathcal{L} \vdash s'$
 - (4) if $\gamma_i :: H; \mathcal{L} \vdash s_i$ for each $i \in 1..n$, then $(\prod_{i \in 1..n} \gamma_i) :: H; \mathcal{L} \vdash \bigwedge_{i \in 1..n} s_i$
 - (5) if $H; \mathcal{L} \vdash s \searrow \gamma$ then $\gamma :: H; \mathcal{L} \vdash s$
 - (6) if $\gamma :: H; \mathcal{L} \vdash s$ then $H; \mathcal{L} \vdash s \searrow \gamma'$ for some γ' .
-

We can now show that compilation preserves the behavior of programs. More precisely, if a program P compiles to a program P' under a system state H , and the programs are run from a system state H' such that $H \leq H'$, then P and P' evaluate to the same state.

Theorem 4.1 (Compilation correctness). *Suppose that Axioms (1–6) hold, and $(\emptyset, \emptyset, \emptyset) \vdash P \xrightarrow{H, \theta, \chi} P'$. Then for all A and $\rho = (H', -, -)$ such that $H \leq H'$, we have $\rho, P \xrightarrow{\theta, A}^* \rho', Q$ for some Q if and only if $\rho, P' \xrightarrow{\theta, A}^* \rho', Q'$ for some Q' . ($\xrightarrow{\theta, A}^*$ denotes the reflexive-transitive closure of $\xrightarrow{\theta, A}$)*

Finally, we show that a compiled program can never fail due to an access check, if the policy does not change between compile time and run time. Formally, compilation preserves the behavior of programs even if the compiled programs are run without access checks.

Definition 4.2 (\Rightarrow). *Let \Rightarrow be the same reduction relation as \rightarrow except that the rule (**Reduct shell**) is replaced by the following rule, which differs from the earlier version in that its premises do not mention any proofs.*

$$\frac{\begin{array}{c} \theta = (\Delta, \mathcal{L}) \\ \Delta(n) = \eta_1, \dots, \eta_k \quad \rho = (H, S, \xi) \quad H, S \xrightarrow{n(t_1, \dots, t_k)} H', S' \quad \rho' = (H', S', \xi) \end{array}}{\rho, \text{shell } n(t_1, \dots, t_k); P \xrightarrow{\theta, \chi} \rho', P}$$

Theorem 4.2 (Access control redundancy). *Suppose that Axioms (1–6) hold, and $(\emptyset, \emptyset, \emptyset) \vdash P \xrightarrow{H, \theta, \chi} P'$. Then for all A and $\rho = (H', -, -)$ such that $H \leq H'$, we have $\rho, P' \xrightarrow{\theta, A}^* \rho', Q$ for some Q if and only if $\rho, P' \xRightarrow{\theta, A}^* \rho', Q'$ for some Q' .*

Before we close this section, let us point out some consequences of our axioms. Axioms (2), (3), (4), (5) represent standard expectations from the proof system and the theorem prover. Axiom (6) is required to prove soundness of the compiler (“if” direction of Theorem 4.1) since, in its absence, there is no guarantee that a statically provable authorization will be successfully proved in the rule (**Reduct assert**) when executing the source program directly. Axiom (1) is needed for a similar purpose; without this axiom, the compiler must throw away assumptions on the system state in the continuation of any shell command. However, the axiom may seem too strong and invalid in practice. Fortunately, weaker versions of this axiom suffice to prove our theorems for specific programs. In particular, the definition of $H \leq H'$ may be qualified to require that $H \models \varphi$ imply $H' \models \varphi$ for only those φ that appear in a program of interest (and their substitution instances).

5 Implementation

We have implemented a prototype PCAL compiler and tested it on the proof-carrying file system PCFS [18]. The specific logic currently used in our implementation is BL [16, 18]. The interested reader can find a complete example (involving homework management between instructors and students of various courses) in the appendix. We now discuss some implementation details that are left abstract in Section 4.

Rewrite rules

A set of rewrite rules over terms, modeling abstract properties of the concrete implementations of function symbols, can be provided to the compiler to improve its precision. The compiler constructs a normalization function based on these rules, and applies this function eagerly to substitutions. This works well even in cases where it is not possible to interpret function symbols with directed clauses in the policy. (Modeling equations as clauses usually causes proof searches to loop.)

Quantified proofs

Statically generated proofs may contain free variables, and as such they are *parametric* over those variables. In the formal semantics, such proofs are bound and carried as values in the language (in `inject` statements), so they get implicitly instantiated before injection at run time. However in our actual implementation, such proofs are output to temporary files with distinct names (under `/.pf`), and the names are carried in the language; so the free variables in such proofs must be explicitly substituted at run time. This explains why we considered an explicit `-subst` option to the `inject` command in Section 2.

6 Conclusion

PCAL combines static checks and dynamic theorem proving to automate correct and efficient use of a PCA-based interface. PCAL's compiler is modular: it is parametric over both the shell commands (system interface) and the logic it supports. Although this makes the compiler flexible, the interaction between the core language, shell commands, and the logic is subtle and requires careful design. The compiler is made practical through a combination of simple user annotations, static constraint tracking, dynamically checked assertions, and run time support from a command line theorem prover. We prove formally that these ideas work well together. It is our belief that PCAL's design is novel, and that it will be a useful stepping stone for languages that support rule-based access control interfaces in future.

There are several interesting avenues for future work. An obvious one is to run realistic examples on PCAL, to determine what other features are needed in practice. Another possible direction is a code execution architecture where a trusted PCAL compiler is used to generate certified scripts that are run with minimal access control checks. Finally, it will be interesting to apply ideas from PCAL, particularly the use of an automatic theorem prover, in the context of language-based security for access control interfaces (*e.g.*, [4, 20]).

References

- [1] Martín Abadi. Access control in a core calculus of dependency. *Electronic Notes in Theoretical Computer Science*, 172:5–31, 2007. *Computation, Meaning, and Logic: Articles dedicated to Gordon Plotkin*.

- [2] Martín Abadi, Michael Burrows, Butler Lampson, and Gordon Plotkin. A calculus for access control in distributed systems. *ACM Transactions on Programming Languages and Systems*, 15(4):706–734, 1993.
- [3] Andrew W. Appel and Edward W. Felten. Proof-carrying authentication. In *ACM Conference on Computer and Communications Security (CCS '09)*, pages 52–62. ACM, 1999.
- [4] Kumar Avijit, Anupam Datta, and Robert Harper. PCML₅: A language for ensuring compliance with access control policies, 2009. Draft, personal communication.
- [5] Lujo Bauer. *Access Control for the Web via Proof-Carrying Authorization*. PhD thesis, Princeton University, 2003.
- [6] Lujo Bauer, Scott Garriss, Jonathan M. McCune, Michael K. Reiter, Jason Rouse, and Peter Rutenbar. Device-enabled authorization in the Grey system. In *Information Security Conference (ISC '05)*, LNCS, pages 431–445, 2005.
- [7] Moritz Y. Becker, Cédric Fournet, and Andrew D. Gordon. Design and semantics of a decentralized authorization language. In *IEEE Computer Security Foundations Symposium (CSF '07)*, pages 3–15. IEEE, 2007.
- [8] Jesper Bengtson, Karthikeyan Bhargavan, Cédric Fournet, Andrew Gordon, and Sergio Maffei. Refinement types for secure implementations. In *IEEE Computer Security Foundations Symposium (CSF '08)*, pages 17–32. IEEE, 2008.
- [9] Avik Chaudhuri and Martín Abadi. Secrecy by typing and file-access control. In *IEEE Computer Security Foundations Workshop (CSFW'06)*, pages 112–123. IEEE, 2006.
- [10] Avik Chaudhuri and Deepak Garg. PCAL: Language support for proof-carrying authorization systems. Technical Report CMU-CS-09-141, Carnegie Mellon University, 2009.
- [11] Avik Chaudhuri, Prasad Naldurg, and Sriram Rajamani. A type system for data-flow integrity on Windows Vista. In *ACM SIGPLAN Workshop on Programming Languages and Analysis for Security (PLAS'08)*, pages 89–100. ACM, 2008.
- [12] John DeTreville. Binder, a logic-based security language. In *IEEE Symposium on Security and Privacy (S&P'02)*, pages 105–113. IEEE, 2002.
- [13] Cormac Flanagan. Hybrid type checking. In *ACM Symposium on Principles of Programming Languages (POPL'06)*, pages 245–256. ACM, 2006.
- [14] Cédric Fournet, Andrew Gordon, and Sergio Maffei. A type discipline for authorization in distributed systems. In *IEEE Computer Security Foundations Symposium (CSF '07)*, pages 31–48. IEEE, 2007.

- [15] Deepak Garg. Principal-centric reasoning in constructive authorization logic. Technical Report CMU-CS-09-120, Carnegie Mellon University, 2009.
- [16] Deepak Garg. Proof search in an authorization logic. Technical Report CMU-CS-09-121, Carnegie Mellon University, 2009.
- [17] Deepak Garg and Frank Pfenning. Non-interference in constructive authorization logic. In *IEEE Computer Security Foundations Workshop (CSFW '06)*, pages 283–293. IEEE, 2006.
- [18] Deepak Garg and Frank Pfenning. A proof-carrying file system. Technical Report CMU-CS-09-123, Carnegie Mellon University, 2009.
- [19] Yuri Gurevich and Itay Neeman. DKAL: Distributed-knowledge authorization language. In *21st IEEE Symposium on Computer Security Foundations (CSF-21)*, 2008.
- [20] Limin Jia, Jeffrey A. Vaughan, Karl Mazurak, Jianzhou Zhao, Luke Zarko, Joseph Schorr, and Steve Zdancewic. Aura: A programming language for authorization and audit. In *ACM International Conference on Functional Programming (ICFP '08)*. ACM, 2008.
- [21] Butler Lampson, Martín Abadi, Michael Burrows, and Edward Wobber. Authentication in distributed systems: Theory and practice. *ACM Transactions on Computer Systems*, 10(4):265–310, November 1992.
- [22] Chris Lesniewski-Laas, Bryan Ford, Jacob Strauss, Robert Morris, and M. Frans Kaashoek. Alpaca: Extensible authorization for distributed services. In *ACM Conference on Computer and Communications Security (CCS '07)*. ACM, 2007.
- [23] Andrew Pimlott and Oleg Kiselyov. Soutei, a logic-based trust-management system. In *Eighth International Symposium on Functional and Logic Programming (FLOPS 2006)*, pages 130–145, 2006.
- [24] Jeffrey A. Vaughan, Limin Jia, Karl Mazurak, and Steve Zdancewic. Evidence-based audit. In *IEEE Symposium on Computer Security Foundations (CSF '08)*. IEEE, 2008.

A Proofs

A.1 Proof of Theorem 4.1

In the following proof, we use a non-deterministic version of our compilation relation, with the following changes.

1. (**Comp static**) can guess any proof that works, not necessarily a proof returned by the theorem prover.

$$\begin{array}{c}
\Gamma = (\sigma, \Phi, \Xi) \quad \theta = (-, \mathcal{L}) \\
\gamma' :: H, \Phi; \mathcal{L} \vdash |\Xi| \Rightarrow \mathbf{auth}(\chi, \eta, t\sigma) \quad \gamma = \gamma' \langle \Xi \rangle = \bar{\gamma}\sigma \\
\Xi' = \Xi[(\chi, \eta, t\sigma) \mapsto \gamma] \quad \Gamma' = (\sigma, \Phi, \Xi') \quad \Gamma' \vdash P \xrightarrow{H, \theta, \chi} P' \\
\hline
\Gamma \vdash \mathbf{assert}(\eta, t); P \xrightarrow{H, \theta, \chi} \mathbf{inject}(\eta, t) \bar{\gamma}; P'
\end{array}$$

2. **(Comp dynamic)** can be applied whenever **(Comp static)** can be applied.

$$\begin{array}{c}
\Gamma = (\sigma, \Phi, \Xi) \quad \theta = (-, \mathcal{L}) \quad \alpha \text{ fresh in } \Gamma, P \\
\Xi' = \Xi[(\chi, \eta, t\sigma) \mapsto \alpha] \quad \Gamma' = (\sigma, \Phi, \Xi') \quad \Gamma' \vdash P \xrightarrow{H, \theta, \chi} P' \\
\hline
\Gamma \vdash \mathbf{assert}(\eta, t); P \xrightarrow{H, \theta, \chi} \alpha = \mathbf{prove}(\eta, t); \mathbf{inject}(\eta, t) \alpha; P'
\end{array}$$

Note that these changes allow some more programs to be compiled, but no less; thus, they do not affect the soundness of our results.

We begin with a few basic lemmas. We assume that if $(\sigma, \Phi, \Xi) \vdash P \xrightarrow{H, \theta, \chi} P'$ then $(\text{fv}(\Phi) \cup \text{fv}(\Xi)) \cap \text{dom}(\sigma) = \emptyset$; note that this invariant is preserved by our typing rules.

Lemma A.1. *Suppose that $(\sigma, \Phi, \Xi) \vdash P \xrightarrow{H, \theta, \chi} P'$ and $x \notin \text{dom}(\sigma)$.*

Then $(\sigma[x \mapsto t\sigma], \Phi\{t\sigma/x\}, \Xi\{t\sigma/x\}) \vdash P \xrightarrow{H, \theta, \chi} P'$.

Proof. By induction on the derivation of the compilation judgment for P . The only interesting case is **(Comp static)**, where we use Axiom (2).

Case $P = \mathbf{assert}(\eta, t); \widehat{Q}$ and $P' = \mathbf{inject}(\eta, t) \gamma; \widehat{Q}'$

such that $\theta = (-, \mathcal{L})$,

$\gamma' :: H, \Phi; \mathcal{L} \vdash |\Xi| \Rightarrow \mathbf{auth}(\chi, \eta, t\sigma)$,

$\gamma = \gamma' \langle \Xi \rangle$,

$\widehat{\Xi} = \Xi[(\chi, \eta, t\sigma) \mapsto \gamma]$,

and $(\sigma, \Phi, \widehat{\Xi}) \vdash \widehat{Q} \xrightarrow{H, \theta, \chi} \widehat{Q}'$.

Let $\sigma' = \{t'\sigma/x\}$.

We require $(\sigma[x \mapsto t'\sigma], \Phi\sigma', \Xi\sigma') \vdash P \xrightarrow{H, \theta, \chi} P'$.

By Axiom (2), $\gamma'\sigma' :: H, \Phi\sigma'; \mathcal{L} \vdash |\Xi\sigma'| \Rightarrow \mathbf{auth}(\chi, \eta, t\sigma\sigma')$.

Also, by assumption, $(\sigma[x \mapsto t'\sigma], \Phi\sigma', \widehat{\Xi}\sigma') \vdash \widehat{Q} \xrightarrow{H, \theta, \chi} \widehat{Q}'$.

The result follows.

□

Lemma A.2. *Suppose that $(\sigma[x \mapsto t\sigma], \Phi, \Xi) \vdash P \xrightarrow{H, \theta, \chi} P'$.*

Then $(\sigma, \Phi, \Xi) \vdash P\{t/x\} \xrightarrow{H, \theta, \chi} P'\{t/x\}$.

Proof. By induction on the derivation of the compilation judgment for P .

□

Lemma A.3. Suppose that $(\sigma, \Phi, \Xi) \vdash P \xrightarrow{H, \theta, \chi} P'$, $\tilde{x} = \mathcal{W}(P)$, and $(\sigma \setminus \tilde{x}, \Phi, \Xi) \vdash Q \xrightarrow{H, \theta, \chi} Q'$. Then $(\sigma, \Phi, \Xi) \vdash P; Q \xrightarrow{H, \theta, \chi} P'; Q'$.

Proof. By induction on the derivation of the compilation judgment for P . The only interesting cases are **(Comp assign)**, where we use Lemma A.1, and **(Comp for)** and **(Comp test)**, which are similar but do not require any additional lemma. Let $\Gamma = (\sigma, \Phi, \Xi)$.

Case $P = \text{for } x \text{ in } t \{ \widehat{P} \}; \widehat{Q}$

Then $P' = \text{for } x \text{ in } t \{ \widehat{P}' \}; \widehat{Q}'$ such that

x is fresh in Γ ,

$\widehat{\Phi} = \Phi, \text{member}(x, t\sigma)$,

$\tilde{z} = \mathcal{W}(\widehat{P})$,

$(\sigma \setminus \tilde{z}, \widehat{\Phi}, \Xi) \vdash \widehat{P} \xrightarrow{H, \theta, \chi} \widehat{P}'$,

and $(\sigma \setminus \tilde{z}, \Phi, \Xi) \vdash \widehat{Q} \xrightarrow{H, \theta, \chi} \widehat{Q}'$.

By assumption, $\tilde{x} = \tilde{z} \cup \mathcal{W}(\widehat{Q})$

and $(\sigma \setminus \tilde{x}, \Phi, \Xi) \vdash Q \xrightarrow{H, \theta, \chi} Q'$.

We need $(\sigma \setminus \tilde{z}, \Phi, \Xi) \vdash \widehat{Q}; Q \xrightarrow{H, \theta, \chi} \widehat{Q}'; Q'$.

The result follows by the inductive hypothesis.

Case $P = \text{test } \varphi \{ \widehat{P} \}; \widehat{Q}$

Then $P' = \text{test } \varphi \{ \widehat{P}' \}; \widehat{Q}'$ such that

$\tilde{z} = \mathcal{W}(\widehat{P})$,

$\widehat{\Phi} = \Phi, \varphi\sigma$,

$(\sigma, \widehat{\Phi}, \Xi) \vdash \widehat{P} \xrightarrow{H, \theta, \chi} \widehat{P}'$,

and $(\sigma \setminus \tilde{z}, \Phi, \Xi) \vdash \widehat{Q} \xrightarrow{H, \theta, \chi} \widehat{Q}'$.

By assumption, $\tilde{z} = \tilde{x} \cup \mathcal{W}(\widehat{Q})$

and $(\sigma \setminus \tilde{x}, \Phi, \Xi) \vdash Q \xrightarrow{H, \theta, \chi} Q'$.

We need $(\sigma \setminus \tilde{z}, \Phi, \Xi) \vdash \widehat{Q}; Q \xrightarrow{H, \theta, \chi} \widehat{Q}'; Q'$.

The result follows by the inductive hypothesis.

Case $P = (x = t; \widehat{Q})$

Then $P' = (x = t; \widehat{Q}')$ such that

$(\sigma[x \mapsto t\sigma], \Phi, \Xi) \vdash \widehat{Q} \xrightarrow{H, \theta, \chi} \widehat{Q}'$.

By assumption, $\tilde{x} = \{x\} \cup \mathcal{W}(\widehat{Q})$

and $(\sigma \setminus \tilde{x}, \Phi, \Xi) \vdash Q \xrightarrow{H, \theta, \chi} Q'$.

We need $(\sigma[x \mapsto t\sigma], \Phi, \Xi) \vdash \widehat{Q}; Q \xrightarrow{H, \theta, \chi} \widehat{Q}'; Q'$.

Let

- $x \in \mathcal{W}(\widehat{Q})$.

The result follows by the inductive hypothesis.

- $x \notin \mathcal{W}(\widehat{Q})$.
Then $\sigma \setminus W(\widehat{Q}) = (\sigma \setminus \tilde{x})[x \mapsto t\sigma]$.
The result follows by Lemma A.1 and the inductive hypothesis.

□

Building on these basic lemmas, the proof of Theorem 4.1 relies on the following main lemma. Here, we write $\bar{\sigma} <_{\mu} \sigma$ when $\bar{\sigma}\mu = \sigma$ and μ only substitutes proofs. Note that since terms t and propositions φ cannot contain proof variables, we always have $t\mu = t$ and $\varphi\mu = \varphi$.

Lemma A.4. *Suppose that Axioms (1–6) hold. Let σ be any ground substitution, and $\bar{\sigma}, \mu$ be such that $\bar{\sigma} <_{\mu} \sigma$. Let $\theta = \neg, \mathcal{L}$. Let \bar{H}, Φ, Ξ, H , and ξ be such that $H \models \bar{H}, \Phi$ and $\forall(\chi, \eta, t) \in \text{dom}(\Xi)$, we have $H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t) \searrow \xi(\chi, \eta, t)$ and $\Xi(\chi, \eta, t)\mu :: H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t)$.*

Suppose that $(\bar{\sigma}, \Phi, \Xi) \vdash P \xrightarrow{\bar{H}, \theta, \chi} P'$.

- If $H, S, \xi, P\sigma \xrightarrow{\theta, \chi} H', S', \xi', Q\sigma'$ for some Q and σ' then
 $H, S, \Xi\mu, P'\sigma \xrightarrow{\theta, \chi} H', S', \Xi'\mu', Q'\sigma'$ for some Q', μ', Φ' , and Ξ' such that
 $\bar{\sigma} <_{\mu'} \sigma'$, $(\bar{\sigma}, \Phi', \Xi') \vdash Q \xrightarrow{\bar{H}, \theta, \chi} Q'$, $H' \models \bar{H}, \Phi'$, and $\forall(\chi, \eta, t) \in \text{dom}(\Xi')$: $H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t) \searrow \xi'(\chi, \eta, t)$ and $\Xi'(\chi, \eta, t)\mu' :: H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t)$.
- If $H, S, \Xi\mu, P'\sigma \xrightarrow{\theta, \chi} H', S', \xi'', Q'\sigma'$ for some ξ'', Q' , and σ' then $\xi'' = \Xi'\mu'$
and $H, S, \xi, P\sigma \xrightarrow{\theta, \chi} H', S', \xi', Q\sigma'$ for some Q, μ', Φ' , and Ξ' such that
 $\bar{\sigma} <_{\mu'} \sigma'$, $(\bar{\sigma}, \Phi', \Xi') \vdash Q \xrightarrow{\bar{H}, \theta, \chi} Q'$, $H' \models \bar{H}, \Phi'$, and $\forall(\chi, \eta, t) \in \text{dom}(\Xi')$: $H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t) \searrow \xi'(\chi, \eta, t)$ and $\Xi'(\chi, \eta, t)\mu' :: H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t)$.
(Note: $\xrightarrow{\theta, \chi}?$ denotes zero or one steps of $\xrightarrow{\theta, \chi}$)

Proof. By induction on the derivation of the compilation judgment for P . Let $\Gamma = (\bar{\sigma}, \Phi, \Xi)$.

Case $P = \text{for } x \text{ in } t \{ \widehat{P} \}; \widehat{Q}$

Then $P' = \text{for } x \text{ in } t \{ \widehat{P}' \}; \widehat{Q}'$ such that

x is fresh in Γ ,

$\widehat{\Phi} = \Phi, \text{member}(x, t\bar{\sigma})$,

$\tilde{x} = \mathcal{W}(\widehat{P})$,

$(\bar{\sigma} \setminus \tilde{x}, \widehat{\Phi}, \Xi) \vdash \widehat{P} \xrightarrow{\bar{H}, \theta, \chi} \widehat{P}'$,

and $(\bar{\sigma} \setminus \tilde{x}, \Phi, \Xi) \vdash \widehat{Q} \xrightarrow{\bar{H}, \theta, \chi} \widehat{Q}'$.

Furthermore, let

- $H' = H$,
 $S' = S$,
 $\xi' = \xi$,

$\sigma' = \sigma$,
 and $Q = \widehat{P}\sigma_1; \dots; \widehat{P}\sigma_k; \widehat{Q}$ such that
 $\text{members}(H, t\sigma) = t_1, \dots, t_k$
 and $\sigma_1 = \{t_1/x\}, \dots, \sigma_k = \{t_k/x\}$.
 Then $Q' = \widehat{P}'\sigma_1; \dots; \widehat{P}'\sigma_k; \widehat{Q}'$.

We require $(\bar{\sigma}, \Phi', \Xi') \vdash Q \xrightarrow{\bar{H}, \theta, \chi} Q'$ for some Φ', Ξ' .

Let $\Phi' = \Phi, \text{member}(t_1, t\bar{\sigma}), \dots, \text{member}(t_k, t\bar{\sigma})$
 and $\Xi' = \Xi$.

Using the freshness assumption on x :

By Lemmas A.1 and A.2 and weakening, $(\bar{\sigma} \setminus \tilde{x}, \Phi', \Xi') \vdash \widehat{P}\sigma_i \xrightarrow{\bar{H}, \theta, \chi} \widehat{P}'\sigma_i$.

Furthermore, by weakening, $(\bar{\sigma} \setminus \tilde{x}, \Phi', \Xi') \vdash \widehat{Q} \xrightarrow{\bar{H}, \theta, \chi} \widehat{Q}'$.

By Lemma A.3, $(\bar{\sigma} \setminus \tilde{x}, \Phi', \Xi') \vdash \widehat{P}\sigma_1; \dots; \widehat{P}\sigma_k; \widehat{Q} \xrightarrow{\bar{H}, \theta, \chi} \widehat{P}'\sigma_1; \dots; \widehat{P}'\sigma_k; \widehat{Q}'$.

By construction, $\text{dom}(\bar{\sigma}') \cap (\text{fv}(\Phi') \cup \text{fv}(\Xi')) = \emptyset$.

The required result follows by Lemma A.1.

Finally, we have

$\bar{\sigma} <_{\mu'} \sigma'$ where $\mu' = \mu$,

$H' \models \bar{H}, \Phi'$ (by assumption on **members** and **member**, since $\bar{\sigma}\mu' = \sigma$),

and $\forall(\chi, \eta, t) \in \Xi' :$

$\Xi'(\chi, \eta, t)\mu' :: H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t)$ and $H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t) \searrow \xi'(\chi, \eta, t)$.

- The converse case is similar.

Case $P = \text{test } \varphi \{ \widehat{P} \}; \widehat{Q}$

Then $P' = \text{test } \varphi \{ \widehat{P}' \}; \widehat{Q}'$ such that

$\tilde{x} = \mathcal{W}(\widehat{P})$,

$\widehat{\Phi} = \Phi, \varphi\bar{\sigma}$,

$(\bar{\sigma}, \widehat{\Phi}, \Xi) \vdash \widehat{P} \xrightarrow{\bar{H}, \theta, \chi} \widehat{P}'$,

and $(\bar{\sigma} \setminus \tilde{x}, \Phi, \Xi) \vdash \widehat{Q} \xrightarrow{\bar{H}, \theta, \chi} \widehat{Q}'$.

Furthermore, let

- $H' = H$,
 $S' = S$,
 $\xi' = \xi$,
 $\sigma' = \sigma$,
 and $Q = \widehat{P}; \widehat{Q}$ such that
 $H \models \varphi\sigma$.
 Then $Q' = \widehat{P}'; \widehat{Q}'$.

We require $(\bar{\sigma}, \Phi', \Xi') \vdash Q \xrightarrow{\bar{H}, \theta, \chi} Q'$ for some Φ', Ξ' .

Let $\Phi' = \widehat{\Phi}$

and $\Xi' = \Xi$.

By Lemma A.3 and weakening, we have the required result.

Finally, we have

$\bar{\sigma} <_{\mu'} \sigma'$ where $\mu' = \mu$,

$H' \models \bar{H}, \Phi'$ (by assumption above),

and $\forall(\chi, \eta, t) \in \Xi'$:

$\Xi'(\chi, \eta, t)\mu' :: H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t)$ and $H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t) \searrow \xi'(\chi, \eta, t)$.

- $H' = H$,
 $S' = S$,
 $\xi' = \xi$,
 $\sigma' = \sigma$,
and $Q = \widehat{Q}$ such that
 $H \not\models \varphi\sigma$.
Then $Q' = \widehat{Q'}$.

We require $(\bar{\sigma}, \Phi', \Xi') \vdash Q \xrightarrow{\bar{H}, \theta, \chi} Q'$ for some Φ', Ξ' .

Let $\Phi' = \Phi$,

and $\Xi' = \Xi$.

By construction, $\text{dom}(\bar{\sigma}') \cap (\text{fv}(\Phi') \cup \text{fv}(\Xi')) = \emptyset$.

So the required result follows by Lemma A.1.

Finally, we have

$\bar{\sigma} <_{\mu'} \sigma'$ where $\mu' = \mu$,

$H' \models \bar{H}, \Phi'$,

and $\forall(\chi, \eta, t) \in \Xi'$:

$\Xi'(\chi, \eta, t)\mu' :: H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t)$ and $H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t) \searrow \xi'(\chi, \eta, t)$.

- The converse case is similar.

Case $P = (x = t; \widehat{Q})$

Then $P' = (x = t; \widehat{Q'})$ such that

$(\bar{\sigma}[x \mapsto t\bar{\sigma}], \Phi, \Xi) \vdash \widehat{Q} \xrightarrow{\bar{H}, \theta, \chi} \widehat{Q'}$.

Furthermore, let

- $H' = H$,
 $S' = S$,
 $\xi' = \xi$,
 $\sigma' = \sigma$,
and $Q = \widehat{Q}\{t\bar{\sigma}/x\}$.
Then $Q' = \widehat{Q'}\{t\bar{\sigma}/x\}$.

We require $(\bar{\sigma}, \Phi', \Xi') \vdash Q \xrightarrow{\bar{H}, \theta, \chi} Q'$ for some Φ', Ξ' .

Let $\Phi' = \Phi$,

and $\Xi' = \Xi$.

By Lemma A.2 we have the required result.

Finally, we have

$\bar{\sigma} <_{\mu'} \sigma'$ where $\mu' = \mu$,

$H' \models \bar{H}, \Phi'$,

and $\forall(\chi, \eta, t) \in \Xi'$:

$\Xi'(\chi, \eta, t)\mu' :: H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t)$ and $H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t) \searrow \xi'(\chi, \eta, t)$.

- The converse case is similar.

Case $P = \text{shell } n(t_1, \dots, t_k); \widehat{Q}$

Then $P' = \text{shell } n(t_1, \dots, t_k); \widehat{Q'}$ such that

$\theta = (\Delta, \mathcal{L})$,

$\Delta(n) = \eta_1, \dots, \eta_k$,

$(\chi, \eta_i, t_i \bar{\sigma}) \in \text{dom}(\Xi)$ for each i ,

and $(\bar{\sigma}, \Phi, \Xi) \vdash \widehat{Q} \xrightarrow{\bar{H}, \theta, \chi} \widehat{Q'}$.

Furthermore, let

- $\xi' = \xi$,
 $\sigma' = \sigma$,
and $Q = \widehat{Q}$ such that
 $H, S \xrightarrow{n(t_1 \sigma, \dots, t_k \sigma)} H', S'$,
 $\xi(\chi, \eta_i, t_i \sigma) = \gamma_i$ for each i ,
and $\gamma_i :: H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta_i, t_i \sigma)$ for each i .
(Note that we have $H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta_i, t_i \sigma) \searrow \xi(\chi, \eta_i, t_i \sigma)$ for each i .
So by Axiom (5) we do not need the proof-checking above.)
Let $\Xi(\chi, \eta_i, t_i \sigma)\mu = \bar{\gamma}_i$ for each i .
By assumption, $\bar{\gamma}_i :: H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta_i, t_i \sigma)$ for each i .
Then $Q' = \widehat{Q'}$.

We require $(\bar{\sigma}, \Phi', \Xi') \vdash Q \xrightarrow{\bar{H}, \theta, \chi} Q'$ for some Φ', Ξ' .

Let $\Phi' = \Phi$,

and $\Xi' = \Xi$.

Then we have the required result.

Finally, we have

$\bar{\sigma} <_{\mu'} \sigma'$ where $\mu' = \mu$,

$H' \models \bar{H}, \Phi'$ (by Axiom (1) and assumption on \leq and \models),

and $\forall(\chi, \eta, t) \in \Xi'$:

$\Xi'(\chi, \eta, t)\mu' :: H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t)$ and $H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t) \searrow \xi'(\chi, \eta, t)$.

- The converse case is similar.

Case $P = \text{assert } (\eta, t); \widehat{Q}$ and $P' = (\alpha = \text{prove } (\eta, t); \text{inject } (\eta, t) \alpha; \widehat{Q'})$

such that $\theta = (-, \mathcal{L})$,

α is fresh,
 $\widehat{\Xi} = \Xi[(\chi, \eta, t\bar{\sigma}) \mapsto \alpha]$,
 $(\bar{\sigma}, \Phi, \widehat{\Xi}) \vdash \widehat{Q} \xrightarrow{H, \theta, \chi} \widehat{Q}'$.
Furthermore, let

- $H' = H$,
 $S' = S$,
 $\xi' = \xi[(\chi, \eta, t\sigma) \mapsto \gamma]$,
 $\sigma' = \sigma[\alpha \mapsto \gamma]$,
and $Q = \widehat{Q}$ such that
 $H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t\sigma) \searrow \gamma$.
Then $Q' = \widehat{Q}'$.

We require $(\bar{\sigma}, \Phi', \Xi') \vdash Q \xrightarrow{\bar{H}, \theta, \chi} Q'$ for some Φ', Ξ' .

Let $\Phi' = \Phi$,
and $\Xi' = \widehat{\Xi}$.

Then we have the required result.

Finally, we have

$\bar{\sigma} <_{\mu'} \sigma'$ where $\mu' = \mu[\alpha \mapsto \gamma]$,

$H' \models \bar{H}, \Phi'$,

and $\forall(\chi, \eta, t) \in \Xi'$:

$\Xi'(\chi, \eta, t)\mu' :: H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t)$ and $H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t) \searrow \xi'(\chi, \eta, t)$

(since in particular, $H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t\sigma) \searrow \xi'(\chi, \eta, t\sigma)$)

and since $\Xi'(\chi, \eta, t\bar{\sigma})\mu' = \alpha\mu' = \gamma = \xi'(\chi, \eta, t\sigma)$,

so, by Axiom (5), $\Xi'(\chi, \eta, t\sigma)\mu' :: H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t)$.

- $H' = H$,
 $S' = S$,
 $\xi'' = \xi$,
 $\sigma' = \sigma$,
and $Q' = \mathbf{inject}(\eta, t) \gamma; \widehat{Q}$ such that
 $H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t\sigma) \searrow \gamma$.
Then $Q = P$.

We require $(\bar{\sigma}, \Phi', \Xi') \vdash Q \xrightarrow{\bar{H}, \theta, \chi} Q'$ for some Φ', Ξ' .

Let $\Phi' = \Phi$,

and $\Xi' = \Xi$.

By assumption, we already have $H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t\sigma) \searrow \gamma$.

Now we require $(\bar{\sigma}, \Phi, \Xi[(\chi, \eta, t\sigma) \mapsto \gamma]) \vdash \widehat{Q} \xrightarrow{\bar{H}, \theta, \chi} \widehat{Q}'$.

Since $\alpha \notin \text{dom}(\bar{\sigma})$, the result follows by Lemma A.1.

Finally, we have

$\bar{\sigma} <_{\mu'} \sigma'$ where $\mu' = \mu$,

$H' \models \bar{H}, \Phi'$,

and $\forall(\chi, \eta, t) \in \Xi'$:

$\Xi'(\chi, \eta, t)\mu' :: H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t)$ and $H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t) \searrow \xi'(\chi, \eta, t)$.

Case $P = \text{assert } (\eta, t); \widehat{Q}$ and $P' = \text{inject } (\eta, t) \bar{\gamma}; \widehat{Q'}$

such that $\theta = (-, \mathcal{L})$,

$\bar{H}, \Phi; \mathcal{L} \vdash |\Xi| \Rightarrow \mathbf{auth}(\chi, \eta, t\bar{\sigma}) \searrow \gamma'$,

$\bar{\gamma} = \gamma' \langle \Xi \rangle$,

$\widehat{\Xi} = \Xi[(\chi, \eta, t\bar{\sigma}) \mapsto \bar{\gamma}]$,

and $(\bar{\sigma}, \Phi, \widehat{\Xi}) \vdash \widehat{Q} \xrightarrow{\bar{H}, \theta, \chi} \widehat{Q'}$.

Furthermore, let

- $H' = H$,
 $S' = S$,
 $\xi' = \xi[(\chi, \eta, t\sigma) \mapsto \gamma]$,
 $\sigma' = \sigma$,
and $Q = \widehat{Q}$ such that
 $H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t\sigma) \searrow \gamma$.
Then $Q' = \widehat{Q'}$
and $\Xi' = \widehat{\Xi}$.

We require $(\bar{\sigma}, \Phi', \Xi') \vdash Q \xrightarrow{\bar{H}, \theta, \chi} Q'$ for some Φ', Ξ' .

Let $\Phi' = \Phi$.

Then we have the required result.

Finally, we have

$\bar{\sigma} <_{\mu'} \sigma'$ where $\mu' = \mu$,

$H' \vdash \bar{H}, \Phi'$,

and $\forall(\chi, \eta, t) \in \Xi'$:

$\Xi'(\chi, \eta, t)\mu' :: H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t)$ and $H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t) \searrow \xi'(\chi, \eta, t)$

(since in particular, $H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t\sigma) \searrow \xi'(\chi, \eta, t\sigma)$, and

since by Axiom (5), $\gamma' :: H; \mathcal{L} \vdash |\Xi| \Rightarrow \mathbf{auth}(\chi, \eta, t\sigma)$,

and by Axiom (4), $\langle \Xi \rangle \mu :: H; \mathcal{L} \vdash |\Xi|$,

so by Axiom (3), $\bar{\gamma}\mu' :: H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t\sigma)$).

- $H' = H$,
 $S' = S$,
 $\xi'' = \Xi\mu[(\chi, \eta, t\sigma) \mapsto \bar{\gamma}\mu]$,
 $\sigma' = \sigma$,
and $Q' = \widehat{Q'}$.
Let $\Xi' = \widehat{\Xi}$.

Now, by assumption, we have $H; \mathcal{L} \vdash |\Xi| \Rightarrow \mathbf{auth}(\chi, \eta, t\sigma) \searrow \gamma'$.

Also, by Axiom (4), we have $\langle \Xi \rangle \mu :: H; \mathcal{L} \vdash |\Xi|$.

So, by Axiom (5), $\gamma' :: H; \mathcal{L} \vdash |\Xi| \Rightarrow \mathbf{auth}(\chi, \eta, t\sigma)$.

So, by Axiom (3), $\bar{\gamma}\mu' :: H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t\sigma)$.

So, by Axiom (6), $H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t\sigma) \searrow \gamma$ for some γ .

Then $Q = \widehat{Q}$
and $\xi' = \xi[(\chi, \eta, t\sigma) \mapsto \gamma]$.
We require $(\bar{\sigma}, \Phi', \Xi') \vdash Q \xrightarrow{\bar{H}, \theta, \chi} Q'$ for some Φ', Ξ' .
Let $\Phi' = \Phi$.
Then we have the required result.
Finally, we have
 $\bar{\sigma} <_{\mu'} \sigma'$ where $\mu' = \mu$,
 $H' \models \bar{H}, \Phi'$,
and $\forall(\chi, \eta, t) \in \Xi'$:
 $\Xi'(\chi, \eta, t)\mu' :: H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t)$ and $H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t) \searrow \xi'(\chi, \eta, t)$
(since in particular, $H'; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t\sigma) \searrow \xi'(\chi, \eta, t\sigma)$, and
 $\bar{\gamma}\mu' :: H; \mathcal{L} \vdash \mathbf{auth}(\chi, \eta, t\sigma)$ as shown above).

□

Lemma A.5. *Lemma A.4 implies Theorem 4.1.*

Proof. By induction on the length of $\xrightarrow{\theta, A^*}$.

□

A.2 Proof of Theorem 4.2

The proof follows by observation of the proof of Lemma A.4. In particular, for the case where P is a shell command, using the relaxed (**Reduct shell**) rule suffices to establish the required invariants.

B Example: Homework for Courses

Consider the following idealized scenario for homework management of various courses at an university. There is a directory, `"/courses"`, containing the directories of all courses. In each course directory, there is a directory named `"instructor"` and a directory named `"students"`, both containing directories named after all students. Furthermore, the `"instructor"` directory contains a file called `"homework"`, and each directory under `"students"` has a file called `"solution"`.

Figure 4 shows a PCAL program that does the following. It navigates into the `"instructor"` directory, and copies the `"homework"` file into each directory under `"students"` in turn. Then, it navigates into those directories and copies the `"solution"` file of that student into the corresponding directory under `"instructor"`.

Next, we show the policy in effect. The policy is written in the authorization logic BL [15, 16]. In order to represent policies made by different principals, BL includes a modality A **says** s which means that administrator A states, or believes formula s (s usually expresses a policy rule). The **says** modality has been considered in prior work (e.g., [1, 2, 17, 21]), but the inference rules defining its meaning vary. BL's rules are shown below. In addition, any complete axiomatization of first-order intuitionistic logic is also assumed.

```

1 courses = "/courses";
2
3 assert (read, courses);
4 for course in courses {
5   assert (read, course);
6   for user in course {
7     test suffix (user, "instructor") {
8       instructor = user;
9       assert (read, instructor);
10      for fileinstr in instructor {
11        test suffix (fileinstr, "homework") {
12          assert (read, fileinstr);
13          students = course/"students";
14          assert (read, students);
15          for studdir in students {
16            hwstud = studdir/"homework";
17            assert (write, hwstud);
18            shell cp (fileinstr, hwstud)
19          }
20        }
21      }
22    };
23
24    test suffix (user, "students") {
25      assert (read, user);
26      for userdir in user {
27        studname = base (userdir);
28        solninstr = course/"instructor"/studname/"solution";
29        solnstud = userdir/"solution";
30        assert read solnstud;
31        assert write solninstr;
32        shell cp (solnstud, solninstr)
33      }
34    }
35  }
36 }

```

Figure 4: PCAL program for homework example

$$\vdash (A \text{ says } (s \Rightarrow t)) \Rightarrow ((A \text{ says } s) \Rightarrow (A \text{ says } t)) \quad (\text{K})$$

$$\frac{\vdash s}{\vdash A \text{ says } s} \quad (\text{N})$$

$$\vdash (A \text{ says } s) \Rightarrow (A' \text{ says } A \text{ says } s) \quad (\text{I})$$

$$\vdash A \text{ says } ((A \text{ says } s) \Rightarrow s) \quad (\text{C})$$

In our specific policy, we assume that S and L are separate authorities. The formula $\mathbf{auth}(A, \eta, t)$ is defined as $S \text{ says } \mathbf{may}(A, \eta, t)$; in other words, S represents the enforcer of the policy. On the other hand, L is a local authority that may certify some formulas that S relies on, for example, the validity of the `"/courses"` directory and the membership of certain principals in special groups for which certain policy rules may apply.

We focus on a detailed modeling of the relationship between the function symbol `/` (that concatenates directory paths with file names to give file paths) and directory membership constraints. This allows the compiler to reason, for example, that if f is in directory d and the suffix of f is x then $f \equiv d/x$. Other rules allow principals in the group `special` to inherit `read` permissions from ancestor directories. Finally, there are rules that are specific to instructors and students, specifying which files in the others' directories they are allowed to read and write.

The policy rules are split into two parts. The first part contains the rules stated by L :

$$\begin{aligned} & \forall c. (S \text{ says } \mathbf{member}(c, \text{"/courses"})) \Rightarrow \mathbf{course}(c) \\ & \mathbf{special}(\text{"User"}) \end{aligned}$$

The next part contains the rules stated by S . These include all the policy rules, plus rules that model equivalences between paths constructed using `/`, `base`, and `suffix`.

$$\begin{aligned} & \forall f. \forall d. \forall x. \mathbf{member}(f, d) \Rightarrow \mathbf{suffix}(f, x) \Rightarrow f \equiv d/x \\ & \forall f. \forall d. \forall x. \forall p. \mathbf{member}(f, d) \Rightarrow \mathbf{suffix}(f, x) \Rightarrow d \equiv p \Rightarrow f \equiv p/x \\ & \forall A. \forall f. \forall p. \forall \eta. f \equiv p \Rightarrow \mathbf{may}(A, \eta, f) \Rightarrow \mathbf{may}(A, \eta, p) \\ & \forall f. \mathbf{suffix}(f, \mathbf{base}(x)) \\ & \forall A. \forall c. \forall d. (L \text{ says } \mathbf{course}(c)) \Rightarrow \mathbf{may}(A, \mathbf{read}, c/\text{"instructor"}) \Rightarrow \\ & \quad \mathbf{may}(A, \mathbf{read}, c/\text{"students"}) \\ & \quad \wedge \mathbf{member}(d, c/\text{"students"}) \Rightarrow \mathbf{may}(A, \mathbf{write}, d/\text{"homework"}) \\ & \forall A. \forall c. \forall x. \forall d. (L \text{ says } \mathbf{course}(c)) \Rightarrow \mathbf{may}(A, \mathbf{read}, c/\text{"students"/}x) \Rightarrow \\ & \quad \mathbf{may}(A, \mathbf{write}, c/\text{"instructor"/}x/\text{"solution"}) \\ & \quad \wedge d \equiv c/\text{"students"/}x \Rightarrow \mathbf{may}(A, \mathbf{read}, d/\text{"solution"}) \\ & \forall A. \mathbf{may}(A, \mathbf{read}, \text{"/courses"}) \\ & \forall A. \forall f. \forall d. (L \text{ says } \mathbf{special}(A)) \Rightarrow \mathbf{member}(f, d) \Rightarrow \\ & \quad \mathbf{may}(A, \mathbf{read}, d) \Rightarrow \mathbf{may}(A, \mathbf{read}, f) \end{aligned}$$

With this policy, the PCAL compiler can eliminate *all* `assert` statements in the program of Figure 4, if the program is run by principal `"User"`.